arXiv:1606.08098v1 [stat.AP] 26 Jun 2016

California Reservoir Drought Sensitivity and Exhaustion Risk using Statistical Graphical Models

Armeen Taeb[a,1], J. T. Reager[b], Michael Turmon[b], and Venkat Chandrasekaran[a]

[a]California Institute of Technology, Pasadena, CA, 91125
[b]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

December 2, 2015

Abstract

The ongoing California drought has highlighted the potential vulnerability of state water management infrastructure to multi-year dry intervals. Due to the high complexity of the network, dynamic storage changes across the California reservoir system have been difficult to model using either conventional statistical or physical approaches. Here, we analyze the interactions of monthly volumes in a network of 55 large California reservoirs, over a period of 136 months from 2004 to 2015, and we develop a latent-variable graphical model of their joint fluctuations. We achieve reliable and tractable modeling of the system because the model structure allows unique recovery of the best-in-class model via convex optimization with control of the number of free parameters. We extract a statewide 'latent' influencing factor which turns out to be highly correlated with both the Palmer Drought Severity Index (PDSI, $\rho \approx 0.86$) and hydroelectric production ($\rho \approx 0.71$). Further, the model allows us to determine system health measures such as exhaustion probability and per-reservoir drought sensitivity. We find that as PDSI approaches -6, there is a probability greater than 50% of simultaneous exhaustion of multiple large reservoirs.

1 Introduction

The state of California depends on a complex water management system to meet wide-ranging water demands across a large, hydrologically diverse domain. As part of this infrastructure, California has constructed 1530 reservoirs, and a collective storage capacity equivalent to a year of mean runoff from California rivers [9]. The purpose of this system is to create water storage capacity and extend seasonal water availability to meet national agricultural, residential, industrial, power generation, and recreational needs.

Major statewide California precipitation deficits during water years 20122015 rivaled the most intense 4-year droughts in the past 1200 years [10]. The drought was punctuated by low snowpack in the Sierra Nevada, declining groundwater storage, fallowed agricultural lands, and lower reservoir storage [1, 7, 12]. How long these conditions will persist is unknown, and recent research has explored whether their origin is ascribable to natural climate variability alone or has been increasingly exacerbated by anthropogenic influence [5, 17, 13]. A strong 20152016 El Nino was expected to bring some relief for drought conditions [11], but the 2016 precipitation through Spring was not enough to substantially decrease statewide drought severity. These events illustrate that there exists some sensitivity of California reservoirs to drought conditions, with resulting implications for

national water and agricultural security. Yet this sensitivity has been difficult to quantify due to the size and complexity of the reservoir network.

Recent research has explored models to determine the potential for future exhaustion of US reservoirs such as Lake Mead (e.g., [2, 3]). Such approaches yield potentially valuable information to water managers, but the underlying approach used to statistically model a single reservoir is much simpler than an approach for modeling large network of hundreds of reservoirs, whose complex management is based on multiple economic and sectoral objectives [12]. The hard-to-quantify influence of human operators and lack of system closure have made the modeling and prediction of reservoir storage behavior using physical equations challenging, for example in hydrology or climate models [14]. Thus, it is difficult to say with any certainty how the complex network of California reservoirs responds hydrologically to drought severity, and at what point exhaustion becomes feasible.

We present a statistical approach to the problem of modeling Californias complex reservoir network. We model the network variables with a statistical graphical model and use the historical system behavior to estimate network interactions by solving a constrained maximum-likelihood estimation problem. The robustness and predictive accuracy of the model is greatly improved by imposing sparsity constraints; the form of these constraints allows application of new convex optimization tools to uniquely solve the resulting estimation problem.

In its general form, the model allows automatic recovery of unobserved latent variables that influence network behavior, as well as introduction of key forcing variables, most significantly, drought severity index. This approach allows for a mathematical simplification of a complex network so that network behavior and response to forcing can be better understood than before. The learned model allows us to make statistically significant statements about the sensitivity of a representative subset of Californias reservoirs to drought severity, and the resulting conditions leading to increased probability of exhaustion in the future.

2 Data

Our primary dataset is monthly averages of reservoir volumes, derived from daily time series of volumes downloaded from the California Data Exchange Center (CDEC). We also used secondary data for some covariates.

2.1 Reservoir Time Series

We began with a list of all 60 California reservoirs having daily data during the period of study (January 2004 – April 2015). From this list, we excluded the five reservoirs with more than half of their values undefined or zero, leaving 55 reservoirs. This list of daily values was inspected using a simple continuity criterion and approximately 50 specific values were removed or corrected. Corrections were possible in six cases because values had misplaced decimal points, but all other detected errors were set to missing values. The most common error modes were missing values that were recorded as zero volume, and a burst of errors in the Lyons reservoir during late October 2014 that seem due to a change in recording method at that time.

The final set of 55 reservoir volume time series spans 4595 days over the 136 months in the study period. It contains two full cycles of California drought (roughly, 2004–2010 and 2010–present). Four California hydrological zones are represented, with 25, 20, 6, and 4 reservoirs in the Sacramento, San Joaquin, Tulare, and North Coast zones, respectively.

The volumetric data was averaged from daily down to monthly values, and its strong annual component was removed with a per-month, per-station average, leaving a volume anomaly with

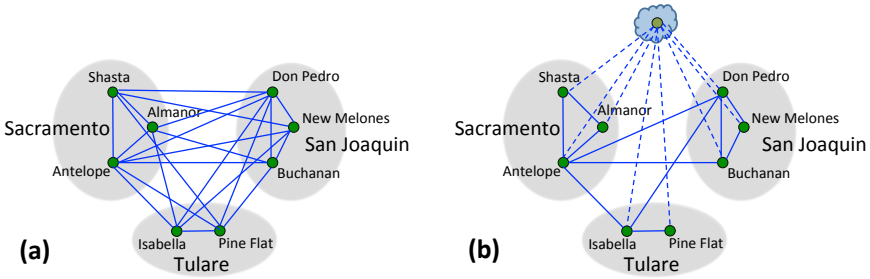

Figure 1: Graphical structure between a collection of reservoirs – without latent variables (a) – with latent variables (b). Green nodes represent reservoirs (variables) and the clouded green node represents latent variables. Solid blue lines represent edges between reservoirs and dotted edges between reservoirs and latent variables. The reservoirs have been grouped according to hydrological zones.

136 monthly observations to be modeled statistically. Before being used in the fitting algorithms of Section 3, each time series is also rescaled by its standard deviation so that each series has unit variance. The volume anomaly, being a de-trended average of small daily fluctuations, is plausibly modeled by a Gaussian frequency law.

2.2 Covariate Time Series

Our model and analysis makes use of ancillary data, i.e., *covariates*, which are observable variables, exogenous to the model, that may typically affect a large fraction of reservoirs. The particular covariates we used are temperature, snowpack, PDSI, number of agricultural workers, hydroelectric power, and consumer price index (CPI). We defined metrics for statewide snowpack (in high-altitude regions) as well as zone-specific snowpack. We apply a time lag of two months to the covariates temperatures, snowpack, and PDSI. As with the reservoir time series, we remove seasonal patterns with a per-month average, and scale each time series to have unit variance.

3 Methods

We introduce three modeling techniques that are employed on the reservoir dataset. We denote the 55 reservoir volumes by $y \in \mathbb{R}^{55}$, and the nine covariates described in Section 2 by $c \in \mathbb{R}^9$. We denote the 136 monthly observations by $\mathcal{D}_n = \{y^{(i)}\}_{i=1}^n \subset \mathbb{R}^{55}$. In what follows, we assume zero-mean variables, or equivalently, that any known mean has been subtracted from the observations — assured in our case by preprocessing that subtracts a climatology from y and c.

3.1 Graphical modeling

A natural framework for modeling a large network of reservoirs is a Gaussian graphical model (defined with respect to a graph) where nodes index reservoirs and edges denote dependencies between reservoirs. The absence of an edge between a pair of reservoirs implies that the corresponding reservoir volumes are independent conditioned on the volumes of the remaining reservoirs in the network. Figure 1(a) is an example of a graphical model over a collection of reservoirs in three hydrological zones: Sacramento, San Joaquin, and Tulare. In comparison to graphical modeling, standard linear regression techniques suffer from over-fitting in this setting where there is a moderate number of observations compared to the number of reservoirs. We demonstrate this point in the supporting information (SI).

The structural properties of a graphical model are encoded in the sparsity pattern of the inverse covariance matrix (the *precision matrix*) over reservoirs. Specifically, consider a Gaussian graphical

model for reservoir volumes y with a covariance matrix Σ. (Both Σ and the corresponding precision matrix $\Theta = \Sigma^{-1}$ must be symmetric, written $\Theta \in \mathbb{S}^{55}$, and positive definite, written $\Theta \succ 0$.) Then, an edge between reservoirs r and r' is present in the graph if and only if $\Theta_{r,r'} \neq 0$, with larger magnitudes indicating stronger interactions. Equivalently, the volumes of reservoirs r and r' are conditionally independent if and only if $\Theta_{r,r'} = 0$. Thus, learning a Gaussian graphical model is equivalent to estimating a *sparse* precision matrix Θ [18, 8].

3.2 Latent variable graphical modeling

A drawback with graphical modeling is that we do not observe all relevant phenomena: some globally-influential quantities may be *hidden* or *latent* variables. Without accounting for these latent variables, there will be confounding dependencies between many reservoirs. For example, water held by a collection of nearby reservoirs might be influenced by a common snowpack variable. Without observing this common variable, all reservoirs in this set would appear to have mutual links, whereas if snowpack is included in the analysis, the common behavior is explained by a link to the snowpack variable. Accounting for latent structure removes these confounding dependencies and leads to sparser interactions between reservoirs. Figure 1(b) illustrates this point. Recent work [4] developed a principled approach, termed latent-variable graphical modeling, to account for latent variables. Fitting a latent-variable graphical model corresponds to representing Θ (the precision matrix of the reservoir volumes) as the sum of a sparse and a low-rank matrix, rather than just a sparse matrix, as in the case of pure graphical modelling. Here the low-rank component of Θ accounts for the effect of the latent variables, and its rank is the number of latent variables. The sparse component of Θ specifies the conditional graphical model structure underlying the distribution of reservoir volumes conditioned on the latent variables.

Based on this decomposition of Θ, a natural model selection method for learning a latent-variable graphical model is via a regularized maximum-likelihood estimator. Specifically, we fit the n observations $\mathcal{D}_n = \{y^{(i)}\}_{i=1}^n \subset \mathbb{R}^{55}$ to a latent-variable graphical model using the estimator:

$$\begin{aligned} (\hat{S}, \hat{L}) = \underset{S,\, L \in \mathbb{S}^{55}}{\arg\min} \quad & -\ell(S - L; \mathcal{D}_n) + \lambda(\gamma \|S\|_1 + \mathrm{tr}(L)) \\ \text{s.t.} \quad & S - L \succ 0,\, L \succeq 0 \quad . \end{aligned} \tag{1}$$

The term $\ell(\Theta; \mathcal{D}_n)$ is the Gaussian log-likelihood of Θ, which, after discarding additive constants and scaling, is the concave function

$$\ell(\Theta; \mathcal{D}_n) = \log\det(\Theta) - \mathrm{tr}\left[\Theta \cdot \tfrac{1}{n} \sum_{i=1}^{n} y^{(i)} y^{(i)'}\right] \quad . \tag{2}$$

The constraints $\succ 0$ and $\succeq 0$ impose positive-definiteness and positive semi-definiteness. Here $\hat{S}$ provides an estimate for the sparse component of Θ (corresponding to the conditional precision matrix of reservoir volumes), and $\hat{L}$ provides an estimate for the low-rank component of Θ (corresponding to the effect of latent variables on reservoir volumes). We note that constraining L to be the zero matrix in (1) provides an estimator for learning a sparse Gaussian graphical model; see [8] for further details. The function $\|\cdot\|_1$ denotes the L_1 norm (element-wise sum of absolute values) that promotes sparsity in the matrix S. This regularizer has successfully been employed in many settings for fitting structured sparse models to high-dimensional data (see the survey [16] and the references therein). The role of the trace penalty on L is to promote low-rank structure [6]. The regularization parameter γ provides a trade-off between the graphical model component and the latent component, and the regularization parameter λ provides overall control of the trade-off

between the fidelity of the model to the data and the complexity of the model. For $\lambda, \gamma \geq 0$, the regularized maximum likelihood estimator of (1) is a convex program with a unique optimum. Further theoretical support for this estimator is presented in [4].

3.3 Latent Variable Structure

Latent variable graphical modeling identifies the effect that underlying latent variables have on the reservoir volumes. For applications, the following approach to associate semantics to these latent variables can aid understanding of the estimated model. Suppose we learn a latent-variable graphical model with k latent variables with the methodology just described. Let $z \in \mathbb{R}^k$ denote the latent variables and $y \in \mathbb{R}^{55}$ denote reservoir volumes; further, partition the joint precision matrix of (y, z) as $\tilde{\Theta} = \begin{pmatrix} \tilde{\Theta}_y & \tilde{\Theta}'_{zy} \\ \tilde{\Theta}_{zy} & \tilde{\Theta}_z \end{pmatrix}$. A natural approximation for the time series of z given observations $\mathcal{D}_n$ is the conditional mean:

$$\tilde{z}^{(i)} = \mathbb{E}[z^{(i)} \mid y^{(i)}] = \tilde{\Theta}_z^{-1} \tilde{\Theta}_{zy} y^{(i)}. \tag{3}$$

If $\tilde{\Theta}_z$ and $\tilde{\Theta}_{zy}$ were explicitly known, the length-n time series $\{\tilde{z}^{(i)}\}_{i=1}^n \subset \mathbb{R}^k$ would provide an estimate of the latent variables given observations $\mathcal{D}_n$. As discussed in [4], the low-rank component in the decomposition of the marginal precision matrix of y is $L = \tilde{\Theta}'_{zy} \tilde{\Theta}_z^{-1} \tilde{\Theta}_{zy}$. However, having an estimate for L does not uniquely identify $\tilde{\Theta}_z^{-1} \tilde{\Theta}_{zy}$. Indeed, for any non-singular $A \in \mathbb{R}^{k \times k}$, one can transform $\tilde{\Theta}_z \to A \tilde{\Theta}_z A'$ and $\tilde{\Theta}_{zy} \to A \tilde{\Theta}_{zy}$ without altering L. In terms of z, these observations imply that for any non-singular A, $\{A^{-1} \tilde{z}^{(i)}\}_{i=1}^n$ is an equivalent realization of the latent variable time series: z is recoverable only up to a nonsingular transformation.

Nevertheless, the structure of the low-rank matrix L places a constraint on the effect of the latent variables z on y. Let $\tilde{Z} \in \mathbb{R}^{n \times k}$ denote a (non-unique) realization of latent variable observations. As we have seen, $\tilde{Z} A'^{-1}$ is an equivalent realization. The key *invariant* is the column-space of $\tilde{Z}$, a k-dimensional linear subspace of R^n, which we term the *latent space*, which can be recovered as follows.

Letting $Y \in \mathbb{R}^{n \times 55}$ denote observations of reservoir volumes, (3) becomes $\tilde{Z} = Y \tilde{\Theta}'_{zy} \tilde{\Theta}_z^{-1}$. Since the column-space of $Y \tilde{\Theta}'_{zy} \tilde{\Theta}_z^{-1}$ is equal to the column-space of YL, one set of basis elements for the latent space is the k left singular vectors of the matrix YL, which can be readily computed. We interpret the underlying latent variables by linking observed covariates to this latent space. Specifically, we evaluate the correlation of a covariate with the latent variables by computing the energy of the projection of the time series of the covariate onto the latent space. Covariates with large correlation represent interpretable components of the underlying latent variables: these covariates have significant overlap with the latent space.

3.4 Conditional latent-variable graphical modeling

We provided an approach for linking a latent variable space defined by L to observed covariates. We can extend our modeling framework to incorporate covariates $x \in \mathbb{R}^q$ that are most correlated with the latent space, where q is the number of covariates included from the overall collection $c \in \mathbb{R}^9$ described in Section 1. Since the covariates x can account for the effect of some of the latent variables in the previous modeling framework, the distribution of y given x may still depend on a few *residual latent variables*. Therefore, we fit a latent-variable graphical model to the conditional distribution of $y \mid x$.

In the SI, we discuss how the convex program of (1) can be modified for fitting a conditional latent-variable graphical model to the augmented data $\mathcal{D}_n^+ = \{(y^{(i)}, x^{(i)})\} \subset \mathbb{R}^{55+q}$. This estimator produces three components: $\hat{\Theta}$, $\hat{L}_y$, and $\hat{S}_y$. The matrix $\hat{\Theta} \in \mathbb{S}^{55+q}$ is an estimate for the joint

precision matrix of (y, x), and in particular, its submatrix $\hat{\Theta}_y$ is an estimate for the conditional precision matrix of the reservoirs given the covariates. The submatrix $\hat{\Theta}_y$ can be decomposed as $\hat{\Theta}_y = \hat{S}_y - \hat{L}_y$. The matrix $\hat{L}_y \in \mathbb{S}^{55}$ is an estimate for the effect of residual latent variables on $y \mid x$, with the rank of $\hat{L}_y$ equal to number of residual latent variables. Finally, $\hat{S}_y \in \mathbb{S}^{55}$ is an estimate of the precision matrix of $y \mid x$ conditioned on these residual latent variables.

4 Results

We now apply the methods described above to the reservoir system. For our experiments, we set aside monthly observations of reservoir volumes and covariates from August 2005 – August 2013 as a training set ($n_{\text{train}} = 96$) and monthly observations from January 2004 – July 2005 and September 2013 – April 2015 as a (disjoint) testing set ($n_{\text{test}} = 40$). Let $\mathcal{D}_{\text{train}}$ and $\mathcal{D}^+_{\text{train}}$, respectively, be the training set of reservoir volumes and that of reservoir volumes augmented with covariate data; $\mathcal{D}_{\text{test}}$ and $\mathcal{D}^+_{\text{test}}$ are the corresponding testing sets. The particular covariates that are included in the augmented datasets $\mathcal{D}^+_{\text{train}}$ and $\mathcal{D}^+_{\text{test}}$ are discussed later.

Recall that, for each variable, we remove seasonal components by subtracting a per-month average as computed on the training set, leaving an anomaly that we model as Gaussian. We scale time series to have unit variance on the training observations, and then apply the same transformation to the testing set. The convex programs associated with each modeling framework are regularized max-det programs and are solved in polynomial time using general-purpose solvers [15].

4.1 Graphical Modeling

We begin by learning a graphical model over the reservoirs in our network. Specifically, we use $\mathcal{D}_{\text{train}}$ as input to the appropriate convex program — this is the program (1) with L constrained to be the zero matrix. We select the single regularization parameter λ (γ is set to 1 without loss of generality) by cross-validation with $\mathcal{D}_{\text{test}}$ to learn a graphical model consisting of 385 edges. Recall from Section 3 that the sparsity pattern of the precision matrix $\hat{\Theta}$ encodes the dependence structure between reservoirs, so that a zero entry in the precision matrix implies that the corresponding reservoir volumes are independent conditioned on the volumes of the remaining reservoirs in the network. We denote the strength of an edge as the normalized magnitude of the corresponding precision matrix entry, that is,

$$s(r, r') = |\Theta_{r,r'}| / (\Theta_{r,r}\Theta_{r',r'})^{1/2} \geq 0 \quad .$$

The lower triangle of Figure 2 shows the dependence relationships between reservoirs in this graphical model. The five strongest edges in this graphical structure are between reservoirs Relief–Main Strawberry, Cherry–Hetch Hetchy, Little Grass Valley–Lake Valley, Almanor–Davis, and Coyote Valley–Warm Spring. The presence of these strong edges is sensible as each such edge is between reservoirs in the same hydrological zone, and 4 of these 5 edges are between pairs of reservoirs fed by the same river basin.

We also observe that the majority of interactions in this graphical model are among reservoirs that have similar volumetric capacities. Specifically, 80% of the edges are between reservoirs that are within a factor of 5 in capacity. As a point of comparison, consider a complete graph with all pairwise interactions. In this graphical structure, 55% of the edges are among reservoirs that are within a factor of 5 in capacity. Thus, graphical modelling removes the dependencies between reservoirs of vastly different capacity. This is expected since reservoirs with substantially different

volumetric capacity are unlikely to have similar variability. In a similar vein, we notice that a significant portion of the edges in the graphical model are between reservoirs at similar elevations. For example, of the ten strongest edges, nine are between reservoirs that are within a factor of 1.7 in altitude.

Finally, we observe that a large portion of the strong interactions occur between reservoirs in the same hydrological zone, here denoted $g(r)$. To quantify this observation, we consider

$$\kappa = \frac{\sum_{r \neq r' \text{ and } g(r)=g(r')} s(r,r')}{\sum_{r \neq r'} s(r,r')} \quad ,$$

the ratio of within-zone edge strength to total edge strength. The model we fit has $\kappa = 0.72$, so 72% of interactions were within-zone. Nevertheless, we notice that there are a number of strong edges between reservoirs that are geographically far apart. This suggests the presence of global state-wide latent variables (Figure 1(b)) that influence distant reservoirs, and accounting for these variables may remove confounding relationships between geographically distant reservoirs.

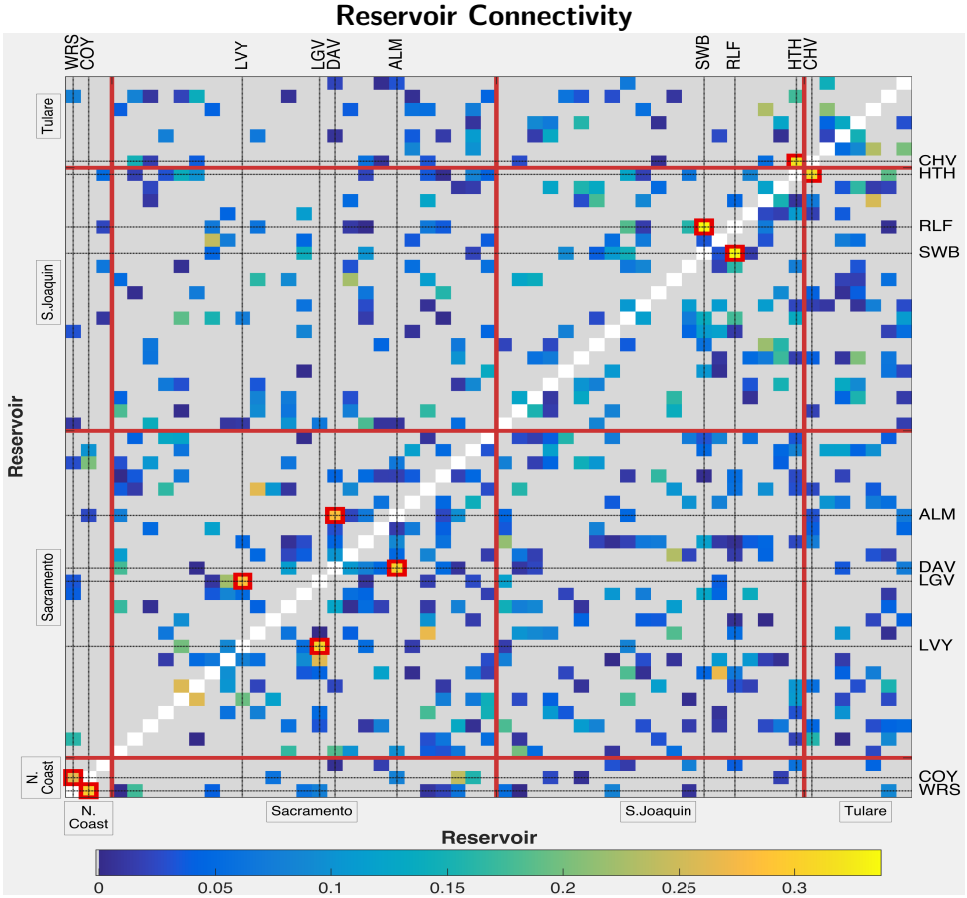

Figure 2: Linkages between reservoir pairs in the latent-variable sparse graphical model (upper triangle, 2 latent variables) compared with those of the ordinary sparse model (lower triangle). Connection strength $s(r,r')$ is shown in the image map, with unlinked reservoir pairs drawn in gray. The four hydrological zones are separated by red lines. Red boxes surround the five strongest connections in each model.

4.2 Latent-Variable Graphical Modeling

We use the method described in Section 3 to account for the effect of latent variables when fitting a graphical model. Specifically, using observations $\mathcal{D}_{\text{train}}$ as input to the convex program (1), and

an appropriate choice of regularization parameters (λ, γ) (chosen via cross-validation with $\mathcal{D}_{\text{test}}$), we learn a latent-variable graphical model with estimates $\hat{S}$ and $\hat{L}$. Recall that $\hat{L}$ represents the effect of the latent variables on reservoir volumes and its rank is the number of latent variables; $\hat{S}$ is the conditional precision matrix of reservoir volumes conditioned on the latent variables, with its sparsity pattern encoding the conditional graphical structure among reservoirs. The best model consists of two latent variables, together with a conditional graphical model (conditioned on the latent variable) having 259 edges.

The upper triangle portion of Figure 2 shows the dependency relationships between reservoir pairs in the sparse component of the latent-variable graphical model. Comparing this graphical structure with the graphical structure without latent variables, we observe that incorporating latent variables weakens or removes many connections between reservoirs: 124 are removed and 232 are weakened by including the latent variables. Of the 124 edges removed, 79 are between reservoirs in different hydrological zones. Further, the ratio of inner zone edge strengths to total edge strength increases from $\kappa = 72\%$ to $\kappa = 79\%$. These results support the idea that latent variables extract global features that are common to all reservoirs, and incorporating them results in more localized interactions.

Finally, of the 55 reservoirs in our system, 35 are used for sourcing hydroelectric power. In the graphical structure without latent variables, there are 155 pairwise edges between reservoirs that are used for generating hydroelectric power. Once the latent variables are incorporated, all but 8 of these edges are weakened or removed from the graphical structure. This suggests that hydroelectric power may be strongly correlated to the latent variables. We verify this hypothesis next.

4.3 Latent-Variable Structure

We associate semantics to the latent variables learned in the latent-variable graphical model. Following the approach presented in Section 3, we correlate each of the 9 covariates described in Section 1 to the two-dimensional latent space. [1] The covariates for Palmer drought severity index (PDSI) and hydroelectric power are dominant with correlation values of $\rho = 0.86$ and $\rho = 0.71$. Secondary covariate influences are due to snowpack, with little influence from temperature. The complete list of correlations is in Table 1. Figure 3 gives a side-by-side comparison of PDSI and the component of PDSI inside the latent space.

Covariate	Correlation
Palmer drought severity index (PDSI)	0.86
Hydroelectric power	0.71
Snow pack: San Joaquin zone	0.46
Snow pack: Tulare zone	0.38
Snow pack: Sacramento zone	0.28
Consumer price index (CPI)	0.24
Number of agricultural workers	0.22
Temperature: San Joaquin basin	0.14
Temperature: Sacramento basin	0.07

Table 1: Covariates and correlations with the latent space

[1] Recall from Section 2 that we apply a time lag of two months to the covariates PDSI, snowpack (all zones), and temperature (all basins).

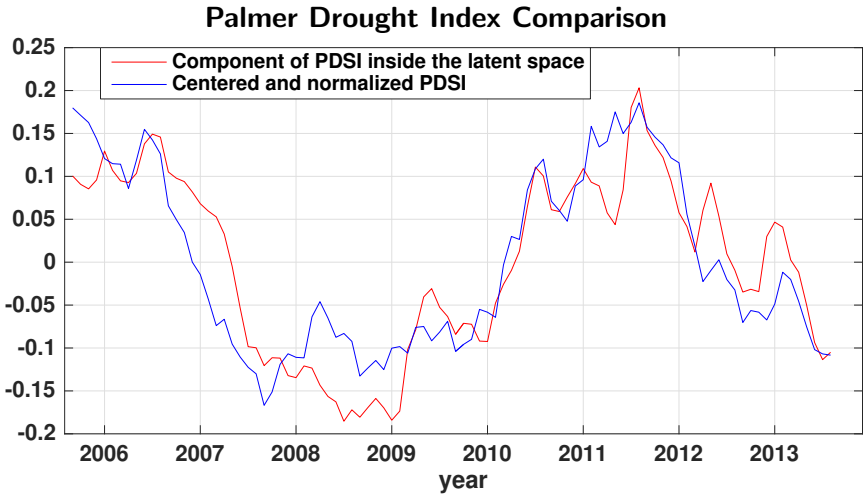

Figure 3: Time series of the component of PDSI inside the latent space compared with PDSI itself. Both time series are centered and scaled to have unit variance.

4.4 Sensitivity to Drought and Exhaustion Probability

Having discovered PDSI as a variable that influences the reservoir system state-wide, we use this covariate to characterize the system-wide response to drought. Our measure for the behavior of a reservoir is *exhaustion probability*: the probability that a reservoir volume drops to zero. To obtain a system-wide response to drought, we compute the probability of exhaustion of a collection of reservoirs conditioned on particular PDSI. We obtain this probability by learning a joint distribution over reservoir volumes and PDSI. We compute this probability for the month of August, when reservoirs are typically at their lowest, but the same calculation applies to any month. Since we applied a time lag of two months to the PDSI time series, these probabilities are computed based on June PDSI.

To learn a joint distribution, let $x \in \mathbb{R}$ denote PDSI and consider a conditional latent-variable graphical model over $(y, x) \in \mathbb{R}^{55+1}$. Using observations $\mathcal{D}^{+}_{\text{train}}$ (consisting of 55 reservoir volumes and PDSI values) as input to the convex program described in section 3.4, we choose parameters (λ, γ) via cross-validation with $\mathcal{D}^{+}_{\text{test}}$ to fit a latent-variable graphical model to the conditional distribution $y \mid x$. The estimated model consists of 1 residual latent variable. Recall that fitting a latent-variable graphical model to the reservoir volumes y (without conditioning on the PDSI covariate) resulted in a model with 2 latent variables. Evidently, by regressing away the effect of PDSI, we are left with one residual latent variable, which supports the observation that PDSI is highly correlated with the latent space in the latent-variable graphical model of y (without PDSI). The conditional latent-variable graphical modeling procedure also provides an estimate of a graphical model of the conditional distribution of y conditioned on x (PDSI) as well as the one residual latent variable – this graphical model consists of 274 edges. Finally, we obtain an estimate $\hat{\Theta} \in \mathbb{S}^{56}$ of the joint precision matrix over (y, x).

Letting $\hat{\Sigma} = \hat{\Theta}^{-1}$, the composite variable $(y, x) \in \mathbb{R}^{56}$ is distributed as $(y, x) \sim \mathcal{N}(0, \hat{\Sigma})$. (Preprocessing to remove climatology causes the mean to be zero.) To determine the behavior of a collection of K reservoirs $\mathbf{r} = \{r_1, r_2, \ldots, r_K\}$ as PDSI varies, we extract the $(K+1) \times (K+1)$ block of $\hat{\Sigma}$ corresponding to $y_{\mathbf{r}} \in \mathbb{R}^K$ and x, and recall that

$$y_{\mathbf{r}} \mid x \sim \mathcal{N}(\hat{\Sigma}_{y_{\mathbf{r}},x}\hat{\Sigma}_x^{-1}x,\ \hat{\Sigma}_{y_{\mathbf{r}}} - \hat{\Sigma}_{y_{\mathbf{r}},x}\hat{\Sigma}_x^{-1}\hat{\Sigma}_{x,y_{\mathbf{r}}}) \quad , \tag{4}$$

an instance of the standard expressions for the conditional mean and variance of these jointly Gaussian variables. Let the August climatology, subtracted during preprocessing, for reservoir

volume y_r ($r \in \mathbf{r}$) be μ_{y_r}, and the June climatology of PDSI x be μ_x. Let the scaling used to make the time series of y_r and x have unit variance be a_{y_r} and a_x. Then, for a June PDSI of u, the probability that at least k of K reservoirs have their volume drop below zero in August is:

$$P(A_K(k) \mid x = a_x(u - \mu_x)) \tag{5}$$

where $A_K(k)$ is the event that $y_r \leq -\mu_{y_r} a_{y_r}$ for at least k of the K reservoirs. The probability in (5), or that of any system-wide event, can be computed using Monte Carlo draws from the joint conditional distribution, (4). Here, we consider those reservoirs having capacity of at least $10^8\,\mathrm{m}^3$ ($K = 31$)[2]. We vary PDSI and compute (5) for selected values of k. Figure 4 indicates a high risk of simultaneous exhaustion for these major California reservoirs.

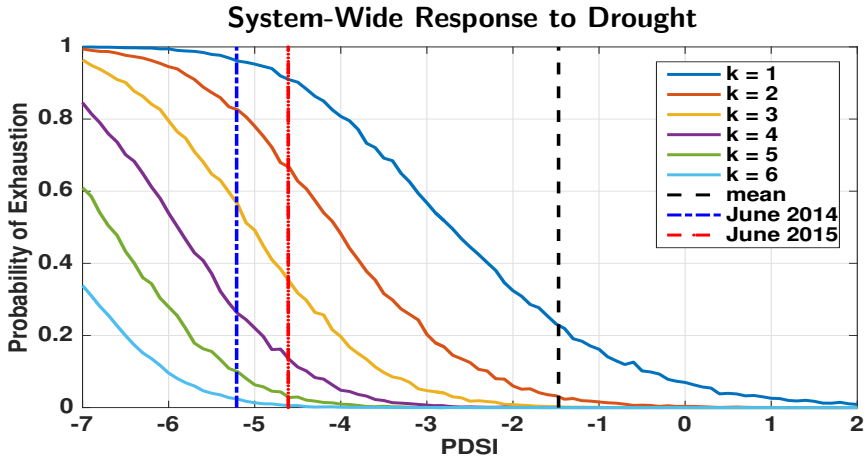

Figure 4: Probability that at least k reservoirs out of 31 *large reservoirs (with capacity* $\geq 10^8\mathrm{m}^3$*) will have volume fall to zero, for a range of PDSI. Dashed black line: average June PDSI (June 2004–June 2015). Dashed blue line: June 2014 PDSI. Dashed red line: June 2015 PDSI.*

[2]Of the 55 reservoirs in our dataset, 22 have capacity below $10^8\,\mathrm{m}^3$. Two of the 33 remaining (Terminus and Success) are flood-control reservoirs: they are unique in that their volume routinely falls below 10% of capacity, independent of PDSI. Thus, we focus on the remaining 31 large reservoirs in what follows.

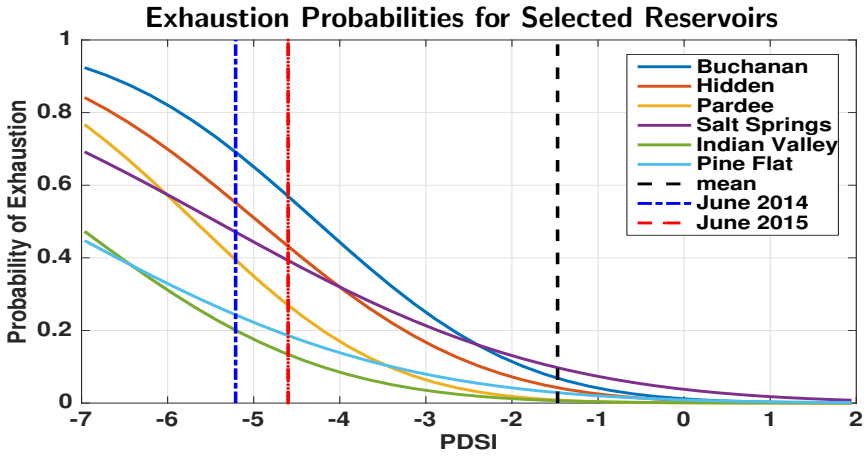

Figure 5: Probability of exhaustion vs. PDSI for 6 large reservoirs (capacity $\geq 10^8 \mathrm{m}^3$) at high risk of exhaustion. Dashed black line: average June PDSI (computed over June 2004-June 2015). Dashed blue line: June 2014 PDSI. Dashed red line: June 2015 PDSI.

To identify particular reservoirs that are at high risk of exhaustion, we compute the exhaustion probability of each reservoir conditioned on PDSI, namely,

$$P(y_r < -\mu_{y_r} a_{y_r} \mid x = a_x(u - \mu_x)) \quad , \tag{6}$$

by applying (4) with $K = 1$. For each reservoir, we sweep over a range of PDSI to compute probabilities of exhaustion. Figure 5 shows those reservoirs (among 31 large reservoirs with capacity greater than $10^8 \mathrm{m}^3$) that were highly sensitive to PDSI. Evidently, these reservoirs are at high risk of exhaustion, and additionally, some have a greater sensitivity to small PDSI changes than others.

5 Discussion

We have used a representative subset of the full California reservoir network (55 of 1530 reservoirs) to model the relationships between reservoirs of different sizes and functions, and the correlation of system response to drought conditions. This unique approach allows us to produce an estimate of the global latent variables and discover that there are highly highly correlated with drought (PDSI) and hydropower generation. We deduce that PDSI, being computed from variables like precipitation and temperature that control mass balance, is a forcing function on system-wide reservoir levels, while correlation of water levels with aggregate hydropower generation is a system-wide response to high reservoir levels across the network. While it is not surprising that PDSI is generally correlated with California reservoir levels, this reduced model allows for the first quantification of system-wide response to drought, a testing of the sensitivity to hypothetical conditions, and a highlighting of the most vulnerable elements of the reservoir network.

These results suggest that with sustained precipitation deficits and a PDSI approaching -6, the probability that four or more of California's major reservoirs (Figure 4) run dry is greater than 50%. This probability increase above 80% as PDSI drops to -7. This does not imply that the reservoirs would in fact run dry. Our approach assumes that the statistical distributions of reservoir response to PDSI changes are stationary, and would exhibit similar behavior under extreme PDSI values to that during the model training period. Realistically, reservoir management behavior will change in response to the most extreme storage conditions. However, the model still demonstrates which reservoirs will be under the greatest threat if typical management behavior were to continue during extreme drought. Additionally, the method used here can forecast other key events that

precede reservoir exhaustion, such as when power generation is made impossible as water levels drop below turbine inlets.

The reservoir system's state is summarized by a complex, dynamic network of correlated time series that respond to a diverse set of global and local drivers, including both natural climate processes and human decision-making. The experimental approach applied here to study reservoirs has the potential to be applicable across many complex data problems because it allows formation and testing of hypotheses about the suspected external drivers on a system, following a traditional experimental strategy. The graphical modeling technique first reduced the number of connections between reservoirs, resulting in a robust model with high predictive power. Then a latent space summarizing all possible configurations of latent-variable times series was estimated by model optimization. Candidate external forcing time series were linked to this latent space to find matches. Once a best match is found (here according to correlation), then the latent variable could be included as a covariate in a new iteration of the graphical modeling procedure and an additional hypothesis for the existence of another latent variable could be tested. This process can be repeated iteratively to identify and represent larger sets of drivers of global system variability. This approach also has value for directing or prioritizing observational efforts.

References

[1] AghaKouchak, A, Cheng, L., Mazdiyasni, O. & Farahmand, A. Global warming and changes in risk and concurrent climate extremes: Insights from the 2014 California drought. *Geophysical Research Letters*. 41:8847–8852, 2014.

[2] Barnett, T. & Pierce, D. When will Lake Mead go dry?. *Water Resources Research*. 44, 2008.

[3] Barnett, T. & Pierce, D. Sustainable water deliveries from the Colorado River in a changing climate. *Proceedings of the National Academy of Sciences*. 106:7334–7338, 2009.

[4] Chandrasekaran, V., Parrilo, P. A. & Willsky, A. S. Latent Variable Graphical Model Selection via Convex Optimization. *Annals of Statistics*, 40:1935–1967, 2012.

[5] Diffenbaugh, N., Swain, D., & Touma, D. Anthropogenic warming has increased drought risk in California. *Proceedings of the National Academy of Sciences*. 112:3931–3936, 2015.

[6] Fazel, M. *Matrix rank minimization with applications*. PhD thesis, Dept. Elec. Engr., Stanford University, 2002.

[7] Famiglietti, J. The global groundwater crisis. *Nature Climate Change*. 4:945–948, 2014.

[8] Friedman, J., Hastie, T. & Tibshirani, R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9:432-441, 2008.

[9] Graf, W. Dam nation: A geographic census of American dams and their large-scale hydrologic impacts. *Water Resources Research*. 35:1305–1311, 1999.

[10] Griffin, D. & Anchukaitis, K. How unusual is the 2012–2014 California drought? *Geophysical Research Letters*. 41:9017–9023, 2014.

[11] Hoell, A., et al. Does El Niño Intensity Matter for California Precipitation? *Geophysical Research Letters*. 43:819–825, 2015.

[12] Howitt, R., Medellín-Azuara, J., MacEwan, D., Lund, J. & Sumner, D. Economic analysis of the 2014 drought for California agriculture. Center for Watershed Sciences, University of California, Davis, 2014.

[13] Shukla, S., Safeeq, M., AghaKouchak, A., Guan, K. & Funck, C. Temperature impacts on the water year 2014 drought in California. *Geophysical Research Letters*.42:4384–4393, 2015.

[14] Solander, K., Reager, J., Thomas, B., David, C. & Famiglietti, J. Simulating human water regulation: The development of an optimal complexity, climate-adaptive reservoir management model for an LSM. *Journal of Hydrometeorology*. 17:725–744, 2016.

[15] Toh, K. C, Todd, M. J. & Tutuncu, R. H. SDPT3 - A MATLAB software package for semidefinite-quadratic-linear programming. Available from `http://www.math.nus.edu.sg/~mattohkc/sdpt3.html`. Accessed April 2016.

[16] Wainwright, M. J. Structured regularizers for high-dimensional problems: Statistical and computational issues. *Annual Review of Statistics and its Applications*. 1:233–253, 2014.

[17] Williams, A., et al. Contribution of anthropogenic warming to California drought during 2012–2014. *Geophysical Research Letters*, 42:6819–6828, 2015.

[18] Yuan, M. & Lin, Y. Model selection and estimation in the Gaussian graphical model. Biometrika, 94:19–35, 2007.

6 Acknowledgement

This research was partly carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with NASA.

7 Supporting Information

7.1 Linear Regression and Comparison with Graphical Modeling

Linear regression is a commonly employed statistical technique for describing the relationship among a collection of variables in a hydrological setting. Applying this framework to our context, each reservoir volume anomaly is modeled as a linear combination of other reservoirs. Thus, a given reservoir volume y_r is modeled as a linear combination

$$y_r = \begin{pmatrix} y_1 \, y_2 \, \cdots \, y_{r-1} & y_{r+1} \, \cdots \, y_{55} \end{pmatrix} \beta_r + \epsilon, \tag{7}$$

where $\beta_r \in \mathbb{R}^{54}$, and the scalar error ϵ is distributed as $\mathcal{N}(0, \sigma^2)$. Given observations $\mathcal{D}_n = \{y^{(i)}\}_{i=1}^n \subset \mathbb{R}^{55}$, the maximum likelihood estimator for the parameter β_r solves the following optimization program:

$$\hat{\beta}_r = \underset{\beta_r \in \mathbb{R}^{54}}{\arg\min} \|Y_r - \tilde{Y}_r \beta_r\|_2^2, \tag{8}$$

where $\|\cdot\|_2$ is the euclidean norm, $Y_r = \left(y_r^{(1)} \, y_r^{(2)} \ldots y_r^{(n)}\right)^T$, and $\tilde{Y}_r = \begin{pmatrix} y_1^{(1)} & y_2^{(1)} & \cdots y_{r-1}^{(1)} & y_{r+1}^{(1)} & \cdots y_{55}^{(1)} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ y_1^{(n)} & y_2^{(n)} & \cdots y_{r-1}^{(n)} & y_{r+1}^{(n)} & \cdots y_{55}^{(n)} \end{pmatrix}$. Given sufficient observations, (8) has a closed-form solution:

$$\hat{\beta}_r = (\tilde{Y}_r^T \tilde{Y}_r)^{-1} \tilde{Y}_r^T Y_r \quad . \tag{9}$$

This procedure can be repeated for every $r \in \{1, 2, \ldots 55\}$ to obtain a complete model.

We now discuss advantages of using graphical modeling over classical linear regression for modeling this reservoir system. We use $\mathcal{D}_n$ as input to (9) to learn a full linear regression model. We compare this model with the graphical model obtained in Section 4. We evaluate the performance of each model via the *coefficient of determination* (R^2) which measures the fraction of variance of each reservoir explained by the model, i.e.:

$$R^2 = 1 - \frac{SS_{res}}{SS_{tot}} \tag{10}$$

For each reservoir y_r, we compute the coefficient of determination specified by our model. Let $\bar{y}_r = \frac{1}{n_{\text{train}}} \sum_{i=1}^{n_{\text{train}}} y_r^{(i)}$ denote the average volume of reservoir y_r over the training set. Furthermore, let $f_r^{(i)}$ denote a model's prediction of reservoir volume y_r at a specified time instance. Then the quantities SS_{tot} and SS_{res} in (10) are given by $SS_{tot} = \sum_{i=1}^{n_{\text{test}}} (y_r^{(i)} - \bar{y}_r)^2$ and $SS_{res} = \sum_{i=1}^{n_{\text{test}}} (f_r^{(i)} - \bar{y}_r)^2$ respectively. We evaluate the quantity $f_r^{(i)}$ by computing the expected volume of reservoir y_r conditioned on the volumes of the remaining reservoirs. For the linear regression model with the estimate $\hat{\beta}_r$, $f_r^{(i)}$ is equal to

$$f_r^{(i)} = \left(y_1^{(i)}, y_2^{(i)}, \ldots y_{r-1}^{(i)}, y_{r+1}^{(i)}, \ldots y_{55}^{(i)}\right) \hat{\beta}_r.$$

Let $\hat{\Theta}$ denote the estimated precision matrix of the graphical model and $\hat{\Sigma} = \hat{\Theta}^{-1}$ the corresponding covariance matrix. Further, let $\hat{\Sigma}_r$ denote the extracted 54×54 covariance matrix among reservoirs $\left(y_1, y_2, \ldots y_{r-1}, y_{r+1}, \ldots y_{55}\right)$, and

$\sigma_r = \left(\Sigma_{y_1,y_r}, \Sigma_{y_2,y_r}, \ldots, \Sigma_{y_{r-1},y_r}, \Sigma_{y_{r+1},y_r}, \ldots, \Sigma_{y_{55},y_r}\right)$ denote a row vector of cross covariances between reservoir y_r and the other 54 reservoirs. Then the quantity $f_r^{(i)}$ corresponding in the graphical model is

$$f_r^{(i)} = \sigma_r \hat{\Sigma}_r^{-1} \left(y_1^{(i)}, y_2^{(i)}, \ldots y_{r-1}^{(i)}, y_{r+1}^{(i)}, \ldots y_{55}^{(i)}\right)^T.$$

Figure 6 compares the R^2 values corresponding to the model obtained via linear regression and the model obtained via graphical modeling. Here, only a subset of reservoirs are plotted as linear regression results in very negative R^2 values for the excluded subset. We observe that for a large portion of the reservoirs, linear regression yields negative R^2 values. This exemplifies the danger of over-fitting. On the other the hand, excluding one reservoir, graphical modeling produces strictly larger R^2 values, indicating that graphical modeling is better suited to modeling this reservoir system.

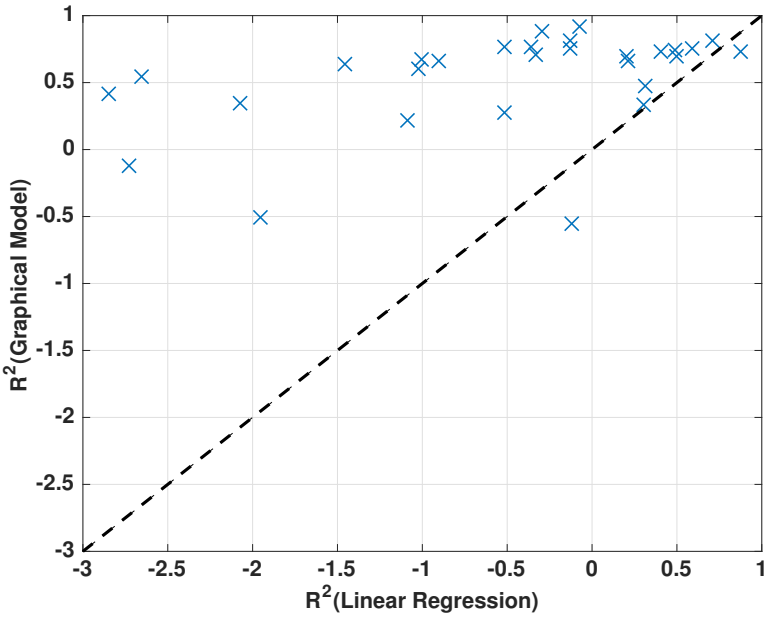

Figure 6: R^2 comparison of a model learned via linear regression and a model learned via graphical modeling. Each cross denotes a reservoir. Crosses above the dotted line suggests that graphical modeling is better suited for modeling the variation of the given reservoir.

7.2 Conditional Latent-Variable Graphical Modeling

In this section, we present an estimator for learning a conditional latent-variable graphical modeling. As described in Section 3, the conditional latent-variable graphical modeling allows for incorporating covariate vectors $x \in \mathbb{R}^q$. Let $\Sigma \in \mathbb{S}^{55+q}$ be the joint covariance matrix over reservoir volumes and covariates $(y, x) \in \mathbb{R}^{55+q}$. A conditional latent-variable graphical model specifies the distribution of $y|x$ as a latent-variable graphical model. Following the description of the latent-variable modeling framework in Section 3, this implies that the conditional precision matrix of y given x must be decomposable as the difference of a sparse and a low-rank matrix. This algebraic property can be naturally stated using the joint precision matrix $\Theta = \Sigma^{-1}$ with conditions on the submatrix Θ_y of $\Theta = \begin{pmatrix} \Theta_y & \Theta_{yx} \\ \Theta_{yx}' & \Theta_x \end{pmatrix}$. The conditional precision matrix of y given x is equal to the submatrix Θ_y; thus Θ_y is the difference of a sparse and low-rank matrix. Based on this algebraic structure desired on

Θ, the following is a natural convex relaxation for fitting the augmented data $\mathcal{D}_n^+ = \{(y^{(i)}, x^{(i)})\} \subset \mathbb{R}^{55+q}$ to this conditional latent-variable graphical model:

$$(\hat{\Theta}, \hat{S}_y, \hat{L}_y) = \underset{\substack{\Theta \in \mathbb{S}^{55+q} \\ S,L \in \mathbb{S}^{55}}}{\arg\min} -\ell(\Theta; \mathcal{D}_n^+) + \lambda(\gamma \|S_y\|_1 + \mathrm{tr}(L_y)) \\ \text{s.t.}\, \Theta \succ 0,\ \Theta_y = S_y - L_y,\ L_y \succeq 0. \tag{11}$$

The term $\ell(\Theta; \mathcal{D}_n^+)$ is the (scaled) Gaussian log-likelihood function over the variables (y, x), given by

$$\ell(\Theta; \mathcal{D}_n^+) = \log\det(\Theta) - \mathrm{tr}\left[\Theta \cdot \tfrac{1}{n} \sum_{i=1}^{n} \begin{pmatrix} y^{(i)} \\ x^{(i)} \end{pmatrix} \begin{pmatrix} y^{(i)} \\ x^{(i)} \end{pmatrix}'\right] \quad , \tag{12}$$

where we have again dropped constant terms and scaled by -2.